

Changes Related To Upgrades

Carefully read this information before upgrading AtScale to the corresponding release.

I2024.1.0

Updating Azure Synapse Data Warehouses

If you use an Azure Synapse data warehouse with an ADLS Gen1 data source, you must migrate to Gen2 and update your data warehouse configuration in AtScale. For more information, see [Adding Azure Synapse Warehouses](#).

I2023.4.1

Upgrading Role-Playing Relationships

I2023.4.1 changes the way role-playing relationships are stored to prevent the creation of ambiguous or duplicative role-playing relationships. You must upgrade your models to use the new role played relationship storage format. Do the following:

1. Open each of your models in Design Center. If any contain duplicate role-playing relationships, a message appears on the **Warnings** tab.
2. Click **Fix It** in the warning message to upgrade the role-playing relationships in the model. Alternatively, you can manually fix the model by deleting all relationships to the affected dimension and recreating them in Design Center (applies to AtScale I2023.4.1 or later).
3. If the model defines semi-additive measures (**Last Non-Empty** or **First Non-Empty**), edit those measures to select the correct semi-additive dimension.

For more information, see [Role-Playing Relationships](#) and [Semi-Additive Measures](#).

ATSCALE-9439

Updating Looker To Use PostgreSQL

If you use Looker 23.16+, you must update your connection to use PostgreSQL after upgrading to AtScale I2023.4.1 or later.



Note: Looker 23.16+ does not support Hive 2. If you use Looker 23.14 or older, you do not need to upgrade your connection.

To update your Looker connection:

1. Update your AtScale settings:

1. In AtScale, go to **Settings > Engine**.
2. Change the value of the **metadataExport** setting from `thrift` to `postgres`.
Note: If you're using the Advanced Settings form, use the `metadataExport.driver.type` key.
3. Click **Save**. You do not need to restart the AtScale engine for the change to take effect.

2. Regenerate your LookML files so they are compatible with PostgreSQL, then copy them to your Looker project. You can either copy the files to Looker manually or publish them via Git. For instructions on both methods, see [Using LookML Files in Looker](#).

Note: If your project contains custom code that uses Hive2 syntax or functions, then you must update your custom expressions with the equivalent PostgreSQL syntax. In particular, be sure to change the column and model attribute delimiters from backticks to quotes.

3. Configure the Looker connection to AtScale:

1. Open Looker and log in.
2. Go to **Admin > Databases > Connections** and click **Add Connection**.
3. Complete the following fields:
 - ▲ **Name:** Enter a name for the AtScale connection.
 - ▲ **Dialect:** Select **PostgreSQL 9.5+**.
 - ▲ **Host:** Enter the name of the database host.
 - ▲ **Port:** Enter `15432`.
 - ▲ **Database:** Enter the name of the published AtScale project.
 - ▲ **Username:** Enter a valid AtScale username.
 - ▲ **Password:** Enter the password for the AtScale user defined by **Username**.



Important: Using port 15432 for the Looker connection may cause the AtScale engine to restart repeatedly. If this occurs, go to **Settings > Organization** in AtScale and change the **Port** setting to `15732`. You must restart the AtScale engine for the change to take effect.

For more information on connecting Looker to AtScale, see [Creating a Looker Connection](#) and [Known Issues and Limitations when Using Looker with AtScale](#).

2023.2.0

Dimensionally Modified Aggregates

After the upgrade from AtScale 2023.1.0 or earlier version to 2023.2.0 or newer version you can easily configure Dimensionally Modified Aggregates (DMA). When you enter a model that contains a Time dimension, a message would be displayed in the Warning Tray in the Design Center, asking you if you wish to enable the default DMA configurations for your Time Hierarchies. If you choose "Yes", the system would add the default DMA configuration for each Level in a Time Hierarchy.

For more information, see [Dimensionally Modified Aggregates](#).

Rebuilding Aggregates

When using SQL Server or Azure Synapse as a data warehouse, after the upgrade from AtScale 2023.1.0 or earlier version to 2023.2.0 or newer version some existing aggregates may be re-built. Please check with AtScale's Support team for additional details and procedure to follow.

2022.4.0

Setting Database For Databricks

When using Databricks as a data warehouse, during the upgrade from AtScale 2022.3.2 or earlier version to 2022.4.0 or newer version you need to do the following:

1. At least a week before the upgrade:
 1. Go to Settings > Organization Settings > Data Warehouses.
 2. Locate the entry for your Databricks data warehouse and choose the Edit button.
 3. Enter the following in the Database field: hive_metastore
 4. Save your changes.
2. Follow the standard procedure to upgrade AtScale; for details, see [Installing and Upgrading AtScale](#).
3. In the Design Center, remap the data sources for your AtScale projects using the corresponding source and database. For details, see [Remapping Data Sources](#).

ATSCALE-9274

2022.3.2

Allowing Dimensionally-Modified Measure Filter Queries

AtScale 2022.3.0 introduced support for Power BI filtering on simple measures; for details, see [Power BI Limitations and Known Issues](#). This changed engine behavior, so that unsupported queries that could return incorrect results fail with the 'Dimensionally-modified measure filters are not supported' message. AtScale 2022.3.2 offers a way to restore the unsupported, pre-2022.3.0 DAX Filter-on-Measure behavior by doing the following:

1. Log in the AtScale as super user.
2. Go to Settings > Organization Settings > Engine.
3. In the Overview section, choose Show Custom Settings.
4. Enter `query.language.dax.blockDimModifiedMeasureFilters` for name, and false for value.
5. Save your changes.

ATSCALE-12068

2022.2.0

Some Projects Will Be Unpublished After Upgrade

Projects with custom Sort keys and undefined Custom Empty Members defaults are now invalid. To eliminate run-time query errors caused by this invalid configuration, affected projects are no longer executable by the AtScale engine.



Published projects with this invalid configuration will not be reachable by query clients after an upgrade. Additionally, it would not be possible to publish such projects in the Design Center.

Before upgrading a production environment to AtScale 2022.2.0, it is recommended that customers first upgrade a development or staging system and inspect the Cube Errors tray in the Design Center to find and fix any missing Custom Empty Member configurations.

It is also possible to make these changes in Design Center with version 2021.4.0. Consider the following:

- ▲ In AtScale 2021.4.0 and earlier versions, the Design Center does not display the affected attributes in the Error Tray.
- ▲ Once fixed, the resulting project can be published and queried.
- ▲ For heavily affected models, downtime can be minimized by performing the work on an upgraded non-production system, and subsequently importing to the production system after an upgrade.

It is recommended that customers schedule extra time to identify and perform the necessary modeling work as part of the upgrade process to AtScale 2022.2.0 or a newer version. For more information, see [Using Custom Empty Members for Levels and Attributes](#).

ATSCALE-7113, ATSCALE-7114

Syntax Validation For Calculated Measures

In earlier releases, when entering the MDX formula for a Calculated measure it was possible to use the equals operator for comparing the CurrentMember function to a scalar value. It was also possible to use the `[Dimension].[Hierarchy].[Level].CurrentMember` syntax. Starting with release 2022.2.0, the syntax validation mechanism detects these cases and provides corresponding warnings.

In case after upgrade to 2022.2.0 you need to keep the existing calculated measures, you can disable the syntax checks introduced in 2022.2.0 by turning on the new `query.language.mdx.currentMember.allowLegacySyntax` engine setting.

ATSCALE-5179

High Cardinality Setting Is No Longer Supported

As described [below](#), the High Cardinality setting in the Design Center was deprecated in AtScale version 2020.2.0. On opening older projects containing this setting users received a corresponding warning.

Starting with AtScale 2022.2.0, the High Cardinality setting is not supported. Also, the Design Center now displays an error instead of a warning.

Before starting the upgrade to AtScale 2022.2.0, it is strongly recommended that customers check the Warning Tray for such projects, click the FIX IT and SAVE buttons, and then republish the projects.

If a project is affected by this deprecation and no action was taken to manually migrate the project, then the AtScale 2022.2.0 installer will make the necessary modifications to the draft project. However, AtScale will take the affected projects off-line after the upgrade. This means that the project must be republished immediately after the upgrade to restore service.



Published projects affected by this deprecation will not be reachable by query clients after an upgrade. Additionally, it would not be possible to publish such projects in the Design Center.

ATSCALE-1868, ATSCALE-9189

Keeping Certificates Imported In Truststore

In releases before 2022.2.0, one of the results from running the `configurator.sh` tool was that the information about certificates in the truststore was deleted, and users needed to import the certificates again. Starting with release 2022.2.0, AtScale can be configured to keep certificates imported in the truststore using the `custom_truststore` option in the `tls` section of the `atscale.yaml` file; for details, see [Configuring TLS](#). If you use TLS, consider using this option when upgrading from 2022.2.0 or newer releases.

2022.1.0

Some Projects Will Be Unpublished After Upgrade

Projects with custom Sort keys and undefined Custom Empty Members defaults are now invalid. To eliminate run-time query errors caused by this invalid configuration, affected projects are no longer executable by the AtScale engine.



Published projects with this invalid configuration will not be reachable by query clients after an upgrade. Additionally, it would not be possible to publish such projects in the Design Center.

Before upgrading a production environment to AtScale 2022.1.0, it is recommended that customers first upgrade a development or staging system and inspect the Cube Errors tray in the Design Center to find and fix any missing Custom Empty Member configurations.

It is also possible to make these changes in Design Center with version 2021.4.0. Consider the following:

- ▲ In AtScale 2021.4.0 and earlier versions, the Design Center does not display the affected attributes in the Error Tray.
- ▲ Once fixed, the resulting project can be published and queried.
- ▲ For heavily affected models, downtime can be minimized by performing the work on an upgraded non-production system, and subsequently importing to the production system after an upgrade.

It is recommended that customers schedule extra time to identify and perform the necessary modeling work as part of the upgrade process to AtScale 2022.1.0 or a newer version. For more information, see [Using Custom Empty Members for Levels and Attributes](#).

ATSCALE-7113, ATSCALE-7114

Configuring Aggregate Maintenance

For new installations of AtScale 2022.1.0 and newer releases, the configuration of the aggregate maintenance job is performed in time-based mode. This means you can have several schedules, each of them specifying a particular day and time when the job should be run.

When upgrading AtScale, consider that the period-based mode for running this job (provided in releases before AtScale 2022.1.0) is no longer available. Now you can configure the job in time-based mode. During upgrade, the aggregate maintenance is set as follows:

- ▲ If it was period-based, it is changed to time-based, using the default schedule (04:00 every day).
- ▲ If it was time-based (for upgrades from version 2021.4.0), it remains as-is.

For more information, see [Configuring Aggregate Maintenance](#).

ATSCALE-6832

2021.4.0

Additional Steps For Upgrading To 2021.4.0 Or Newer Version

When upgrading AtScale from 2021.3.x or earlier version to 2021.4.x or newer you must perform some additional database upgrade steps. For more information, see [Upgrading from AtScale 2021.3 or earlier](#) (standalone) and [Upgrading from AtScale 2021.3 or earlier](#) (cluster).

ATSCALE-5968

2020.2.0

High Cardinality Setting Deprecation

The **High Cardinality** setting present on the Create a Level and Secondary Attribute dialogs were deprecated in AtScale version 2020.2.0. The setting was replaced by two new settings on the same dialogs, Exclude from System-Generated Dimension-Only Aggregates and Exclude from System-Generated Fact-Based Aggregates. These settings allow for the manual exclusion of specific attributes from being used as grouping keys in aggregates.

If upgrading to AtScale version 2020.2.0 the upgrade procedure commences when a data architect opens the Level or Secondary Attribute dialog.

For existing levels marked as **High Cardinality** the migration will do the following:

- ▲ Check **Exclude from System-Generated Dimension-Only Aggregates** for that level or secondary attribute.
- ▲ Check **Exclude from System-Generated Fact-Based Aggregates** for that level or secondary attribute.

If upgrading to AtScale version 2020.2.0 and the upgrade procedure above does not happen, existing dimensional aggregates with the **High Cardinality** setting enabled:

- ▲ Will not be considered for aggregation.

Although the AtScale engine is backwards-compatible with the **High Cardinality** setting, backwards-compatibility will be removed in a future release. It is strongly recommended that you use the Level or Secondary Attribute dialogs to migrate existing **High Cardinality** settings to the new Aggregate Exclusion settings.

ENGINE-4904

2019.3.0

Increased AtScale Metadata Store Write Ahead Log (WAL) File Retention Limit

The AtScale Metadata Store Write Ahead Log (WAL) File Retention limit was increased in AtScale version 2019.3, to increase the recovery window during partial cluster failure situations. The amount of time that the AtScale Administrator has to bring a failed instance of the AtScale Metadata store back online depends on the query load. Under typical load the limit is expected to be roughly 48 hours. This change results in a minimum disk space requirement of 150 GB.

INSTALL-655

2019.1.0

Release Numbering

Starting with the 2019.1.0 release, AtScale uses the following release numbering system: <Scheduled Year>. <Scheduled Quarter>.<Patch Version>

7.4.0

HDP 3.X Requires Additional Data Warehouse Preparation

AtScale requires that Hive and SparkSQL read from the same tables in a single metastore. Additionally SparkSQL cannot read Hive ACID tables. HDP 3.0 changed the default Hive table creation method to ACID and change SparkSQL to read from its own metastore. See [Preparing the Data Warehouse](#) for more information on how to prepare a HDP 3.x hadoop cluster for use by AtScale.

Automatic Rebuild Of User-Defined Aggregates (UDAs) With Exact Count Distinct Measures And At Least One Dimensional Value

UDAs with exact count distinct measures and at least one dimensional value will be automatically rebuilt after upgrade to allow them to be used by queries with the new default grouping by member key values.

Trigger Automatic Aggregate Rebuilds By Republishing All Cubes After Upgrade, Perform Full Cube Rebuilds

After upgrade, some cubes may require rebuilds of old aggregates because of updated semantic information. Cube publication after upgrading is recommended. Expect some automated rebuilding of aggregate tables after publish.

Full cube rebuilds are also recommended.